



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Neural network-based pore flow field prediction in porous media using super resolution

Xu-Hui Zhou, James E. McClure, Cheng Chen, and Heng Xiao

Phys. Rev. Fluids **7**, 074302 — Published 12 July 2022

DOI: [10.1103/PhysRevFluids.7.074302](https://doi.org/10.1103/PhysRevFluids.7.074302)

Neural network based pore flow field prediction in porous media using super resolution

Xu-Hui Zhou,¹ James E. McClure,² Cheng Chen,^{3,*} and Heng Xiao^{1,†}

¹*Kevin T. Crofton Department of Aerospace and Ocean Engineering, Virginia Tech, Blacksburg, Virginia, USA*

²*National Security Institute, Virginia Tech, Blacksburg, Virginia, USA*

³*Department of Civil, Environmental and Ocean Engineering,
Stevens Institute of Technology, Hoboken, New Jersey, USA*

Direct pore-scale simulations of fluid flow through porous media are computationally expensive to perform for realistic systems. Previous works have demonstrated using the geometry of the microstructure of porous media to predict the velocity fields therein based on neural networks. However, such trained neural networks do not perform well for unseen porous media with a large degree of heterogeneity. In this study we propose that incorporating a coarse velocity field in the input of neural networks is an effective way to improve the prediction performance. The coarse velocity field can be simulated with a low computational cost and provides global information to regularize the ill-posedness of the learning problem, which is usually caused by the use of local geometries due to the computational resource constraints. We show that incorporating the coarse-mesh velocity field significantly improves the prediction accuracy of the fine-mesh velocity field by comparison to the prediction that relies on geometric information alone, especially for the porous medium with a large interior vuggy pore space. We also show the flexibility of training the network in using coarse velocity fields with various resolutions. The results suggest that even using coarse velocity field with a very low resolution, the predictions are still enhanced and close to the ground truths. The feasibility of the method is further demonstrated by testing the trained network on real rocks. This study highlights the merits of incorporating a coarse-mesh velocity field into the input for neural networks, which provides global, physics-based information for the model, thereby improving the model's generalization capability.

I. INTRODUCTION

Fluid flow through porous media is a pervasive physical phenomenon, which has critical implications and practical applications in a wide range of natural and engineered processes, such as geological carbon storage [1], oil and gas extraction [2], and contaminant management in groundwater [3, 4]. There are many different methods to simulate velocity fields within the porous media. Among these, the most powerful and attractive methods are direct numerical simulation (DNS) methods, which solve the Navier-Stokes equations and provide the velocity fields with highest accuracy at the pore scale. DNS can be implemented using the lattice Boltzmann method (LBM) [5–7], finite volume method [8, 9] or finite element method [10, 11], among others. However, these methods are highly demanding in terms of computational time and memory requirements.

* chen08@vt.edu

† hengxiao@vt.edu

In the past decade, machine learning has emerged as a promising tool in science and engineering community with applications to turbulent flow modeling [12, 13], molecular dynamics [14, 15], and protein structure prediction [16], among others. In various machine learning techniques, deep learning has been particularly successful because of its strong expressive power in the form of deep neural networks. It is capable of learning the complex functional relationship between a set of input and output data that are obtained through experiments or numerical simulations and then making a rapid and cost-effective prediction.

In the modeling of fluid flow through porous media, deep learning has been used to predict permeability or detailed velocity field based on the geometry of solid microstructure. For example, Srisutthiyakorn [17] used both fully connected neural network and convolutional neural network (CNN) to predict the permeabilities of two-dimensional (2D) and three-dimensional (3D) porous media based on their geometric images, showing that the input features at a larger scale could help the neural network capture global pore connections and improve the prediction capability. Wu *et al.* [18] used CNN to predict permeability based on binary images. They found that including porosity and specific surface area in the fully connected layer of the network could dramatically enhance the prediction accuracy. Their choice of informing the CNN by including the porosity (ratio of void space volume to total sample volume) and specific surface area (solid-void interface area per unit sample volume) were inspired by Kozeny–Carman equation [19, 20]. Both works have achieved successes in predicting permeability based on images of microstructure geometries. However, in some applications it is desirable to predict the velocity fields through porous media. This is much more challenging as the output field resides in a very high dimensional space. Wang *et al.* [21] developed a gated U-net model to map geometric images to velocity fields in both 2D and 3D simulation domains. They found that the prediction performance was highly dependent on the complexity of the pore structures. In particular, the trained model did not perform well in testing porous media with more complex pore structures and often had low prediction accuracies. In contrast, the prediction of permeability was much less susceptible to the complexity of the pore structures, even if it was calculated based on the predicted velocities.

The studies reviewed above focused on using neural network to predict the pore flow in either 2D images or 3D domains with small dimensions. In some cases, we may need to calculate or predict the pore-scale velocity field in a large 3D domain. Training the neural network that maps the pore structure to the velocity field over the entire domain is highly demanding in terms of the GPU memory. This is because a sufficiently complex network with a large amount of parameters is needed to describe the mapping in a very high dimensional space. These parameters, as well as their gradients, must be stored locally in the GPU. Furthermore, even a single pair of input and output over the entire domain consumes a considerable amount of GPU memory. As such, people usually segment the entire domain into small subvolumes and use the local geometry as the input for the neural network to predict the pore flow. However, the velocity field is intrinsically a global physics. This is evident from the fact that the velocity field in Darcy flows is proportional to the local pressure gradient, while the pressure is governed by a Poisson equation. This is an elliptic PDE, for which the solution at any point is determined by the entire domain and the boundary conditions, if any. As a result, the problem of predicting velocity field from local geometry of the microstructure is an intrinsically ill-posed problem. Additional input that incorporates global physics must be provided to regularize the ill-posedness of the problem. Santos *et al.* [22] incorporated such an additional input, the time of flight (ToF), for their CNN-based pore flow prediction model. They showed encouraging results by including global information in the input geometry features of their model. The ToF at any point was defined as the shortest viable fluid flow path

to anywhere on the boundaries of the global domain. While this quantity is straightforward to compute and does
 50 contain global information, it is purely geometric information and does not account for the dynamics of the fluid flow.
 This is evident from the fact that the fluid viscosity is not needed in computing the ToF. On the other hand, while
 full-resolved DNS is often extremely expensive, under-resolved simulation on a coarsened mesh (e.g., by a factor of 4
 or 8 in each dimension) can be performed at a fraction of the cost but still provides essential global information of the
 flow, at least for Darcy flows concerned in this work. Mathematically, such global information is obtained by solving
 55 the pressure Poisson equation on a coarsened mesh (or equivalently the time stepping in explicit methods). This
 strategy draws analogy to the widely used multi-grid methods for solving linear equation systems in computational
 physics [23], where the solution error is reduced on a series of successively coarsened meshes.

In view of the important global information contained in, and the low cost to obtain, the coarse velocity field, it
 is natural to incorporate the coarse velocity field in the input to predict or reconstruct the fine-scale velocity field.
 60 In computer graphics, the technique of increasing the resolution of an image from low to high is referred to as super
 resolution [24]. The CNN-based super resolution has been developed and effectively used to reconstruct fine-scale
 turbulent flow from coarse velocities [25–27]. In the studies of porous media, super resolution techniques have shown
 promise as a way to enhance images of real porous media [28]. Additionally, fine-scale velocity field is usually filtered
 and then employed as the coarse-mesh velocity field to aid in the training of neural networks. For example, Santos
 65 *et al.* [29] trained a multiscale neural network using both geometries and velocities with varying scales (i.e., several
 resolutions from coarse to fine). Instead of including these coarsened velocities into the input, they employed them
 as targets in the loss function. They showed encouraging prediction results of permeabilities even for unseen porous
 media with a large degree of heterogeneity. However, lacking the coarsened velocities in the input may render the
 trained model incapable of accurately predicting the fine-scale velocity field in a new domain, which is a far more
 70 difficult task than predicting the permeability alone. As such, we assist the U-net, a successive encoder-decoder
 network with skip connections in between, with the super resolution technique and use both geometry and coarse
 velocities as the model input. The U-net is utilized to extract the geometry features of local pore structures at
 different scales; the super resolution introduces coarse global physics to improve the prediction and also to regularize
 the ill-posedness of this learning problem caused by the use of local geometry.

75 The rest of the paper is structured as follows. Section II describes the procedures to generate labeled data for training
 and testing, as well as the design of the proposed neural network. Section III highlights the predictive capability of
 the proposed network, including superiority of having the coarse velocities in the input, flexibility of input data and
 viability for real-world applications. Finally, Section IV summarizes the paper and points out directions for future
 research.

80 II. METHODOLOGY

We propose using a U-net based convolutional neural network that maps the geometry of microstructure and a
 coarse velocity field to a fully-resolved velocity field. The network is trained by using data consisting of geometry
 representations and their corresponding velocity fields. To this end, we employ artificial porous media generated
 by using randomly distributed spheres with specified statistics (e.g., number and diameter distributions). Then the
 85 raw geometries (i.e., binary images) are processed to obtain Euclidean distance maps (i.e., distance of a point to

the nearest solid surface) in order to incorporate non-local information. The fully-resolved and coarse velocity fields for each pore geometry are obtained by performing DNS using the lattice Boltzmann method (LBM) on different meshes. Each pore geometry, as well as the corresponding velocity field, is then segmented into small subvolumes for training feasibility and efficiency. The labeled data, consisting of pore geometries and coarse velocities and fine-scale velocities of these subvolumes, are subsequently used to train the neural network. In this section, we first describe the generation of the data for training and testing (Section A), including the generation of pore geometries, numerical simulations performed thereon, and segmentation to training subvolumes. We then introduce the architecture of the neural network proposed in this work (Section B).

A. Data generation

We consider using the 3D digital porous media packed with spheres as the training and testing samples. Two types of porous media with distinct inner pore structures are generated: (1) fully-packed porous media and (2) porous media with a large interior vuggy pore space, which are shown in Fig. 1(a), respectively. The second porous media are highly relevant to hydrocarbon energy recovery in carbonate reservoirs, because carbonate rocks are rich in large fractures and vuggy pores, leading to challenges in direct, coupled simulations of Stokes flow and Darcy flow in the vuggy pore space and rock matrix, respectively.

The generation of porous media is implemented using the open-source SpherePackTools package [30], which can generate a porous medium packed with spheres having lognormally-distributed radii and provides control over the final porosity of the porous medium. The domain length of generated porous media in each direction is $L_x = L_y = L_z = 1.0$. In each porous medium, the number of spheres and variance of sphere radius are set as $N_s = 1000$ and $\sigma^2 = 0.1$, respectively. The target porosity, ϕ_{target} , is varied within the range of $[0.25, 0.35]$ to generate five different porous media for training. Three of them require the removal of spheres from the center to create the large vuggy pore space, while the other two remain unchanged. Specifically, we assume a spherical pore space with a radius of 0.15 in the center, where the spheres are removed to obtain the desired porous medium with a large degree of heterogeneity. Similarly, we generate another two porous media as the testing samples: one with $\phi_{\text{target}} = 0.4$, which is out of the range for training porous media, and the other with a large vuggy pore space in the center, which is obtained based on the previous one.

While binary images of porous media do provide a complete description of the geometric microstructure, they are not sufficient as the input to the neural network. Therefore, we preprocess the binary image to obtain a Euclidean distance map that serves as the geometry representation in the input. The Euclidean distance map provides a considerable non-local (but not global) representation for different geometry shapes. Euclidean distance at a point in the pore space, as suggested by its name, is defined as its minimum distance from this point to the nearest boundary of any solid surface. For simplicity the Euclidean distance for any point inside the solid grain is defined as zero. Specifically, the definition is as follows:

$$\mathcal{D}(\mathbf{x}) = \begin{cases} 0, & \text{if } \mathbf{x} \in \Omega \\ \min_{\mathbf{x}' \in \Omega^b} |\mathbf{x} - \mathbf{x}'|, & \text{if } \mathbf{x} \notin \Omega, \end{cases} \quad (1)$$

where $\mathbf{x} (\in \mathbb{R}^3)$ represents the coordinates of a point within the porous medium domain, $|\cdot|$ denotes the Euclidean

norm, Ω and Ω^b denote the space of spheres and boundaries, respectively. The Euclidean distance transform from binary images is implemented using the open-source Python library SciPy [31]. The Euclidean distance maps for two different pore structures (with and without large vuggy pore space) are shown in Fig. 1(b).

In this work, we consider a single-phase creeping flow through the porous media, where the flow velocity is slow while the viscosity is large [32, 33]. The pore-scale flow velocity simulation is performed using the LBM, which is based on microscopic physical models and mesoscale kinetic equations [34, 35]. Compared with conventional fluid-dynamic models, the LBM has many advantages. For example, it is explicit in the evolution equation, simple to implement, natural to parallelize [36–38], and easy to incorporate new physics such as interactions at the fluid-solid interfaces [39, 40]. It is well known that the LBM can recover the incompressible N-S equations [37, 41]:

$$\begin{aligned} \nabla \cdot \mathbf{u} &= 0, \\ \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} - \nu \nabla^2 \mathbf{u} &= -\frac{1}{\rho} \nabla p + \mathbf{g}, \end{aligned} \quad (2)$$

where \mathbf{u} is the pore flow velocity (m/s), ρ is the fluid density (kg/m³), p is the fluid pressure (Pa), ν is the kinematic viscosity (m²/s), and \mathbf{g} is the body acceleration (m/s²). At the Darcy scale which treats the solid grains and pore space as a continuum, the Darcy velocity, \mathbf{U} , is calculated as the flow rate per unit total porous medium cross section area. The flow rate is calculated by integrating the pore-scale velocity field, \mathbf{u} , over the pore space area at the inlet or outlet of the porous medium [?]. The Darcy velocity and the pressure gradient and body force are related using the Darcy’s law:

$$\mathbf{U} = -\frac{k}{\mu} (\nabla p + \rho \mathbf{g}), \quad (3)$$

where \mathbf{U} is the Darcy velocity (m/s), k is the porous medium permeability (m²), and μ is the fluid dynamic viscosity (Pa s). In practice, both the pressure gradient, ∇p , and body force, $\rho \mathbf{g}$, the two terms in the parentheses of Eq. (3), are able to drive fluid flow through the pore space. In other words, these two terms are equivalent in the macroscopic effect for driving fluid flow through the pore space, as illustrated by Eq. (3). In this study, we adopt the common practice of using the body force term to drive pore fluid flow in order to avoid the compressibility issue of the LBM under high pressure gradients [42–45]. Typically, we match the LB model system with the real physics system by matching the Reynolds number (Re), defined as $Re = UD/\nu$, where U is the magnitude of the Darcy flow velocity (m/s) and D is the characteristic length (m) of the porous medium, which is defined as the average grain diameter.

Specifically, we employ the multi-relaxation-time (MRT) LBM to simulate the steady-state pore flow velocity, which is implemented through the open-source lattice Boltzmann methods for Porous Media (LBPM) software package [46, 47]. We apply the periodic boundary condition for fluid flow in all the three principle directions and apply the body force only in the z direction (i.e., $\mathbf{g} = [0, 0, g_z]^T$). The LBM-simulated pore flow velocity is then normalized by the characteristic velocity, u_c , defined as $u_c = \nu/D$, to obtain the dimensionless pore flow velocity, u^* , which is used as the labeled target for training the neural network.

As mentioned above, in this work we use the coarse velocity field as an additional input for the neural network to predict the fine-scale velocity field. To be specific, we coarsen the fine mesh (512^3) by a factor of four in each dimension to obtain the coarse mesh (128^3) and perform the LB simulations in both meshes. To relate the LB model systems between different meshes, we need to match the dimensionless pore flow velocity, u^* . As such, we match the dimensionless driving force, defined as $Fc = g_z D^3/\nu^2$, between the fine and coarse meshes. Particularly, we set g_z

as 10^{-4} and 6.4×10^{-3} (in lattice unit dx/dt^2) in the fine and coarse meshes, respectively. Note that we scale up
 155 the value of g_z in the coarse mesh by a factor of 64 (i.e., 4^3) due to the fact that the characteristic length in the
 fine mesh is four times of that in the coarse mesh in the LB system. In principle, when we match the dimensionless
 driving force, F_c , we match the dimensionless velocity field, u^* , between the coarse and fine meshes. In this study,
 we focus on predicting the dimensionless velocities in the z direction (i.e., u_z^*), but this method can also be applied
 to predicting velocities in the x and y directions.

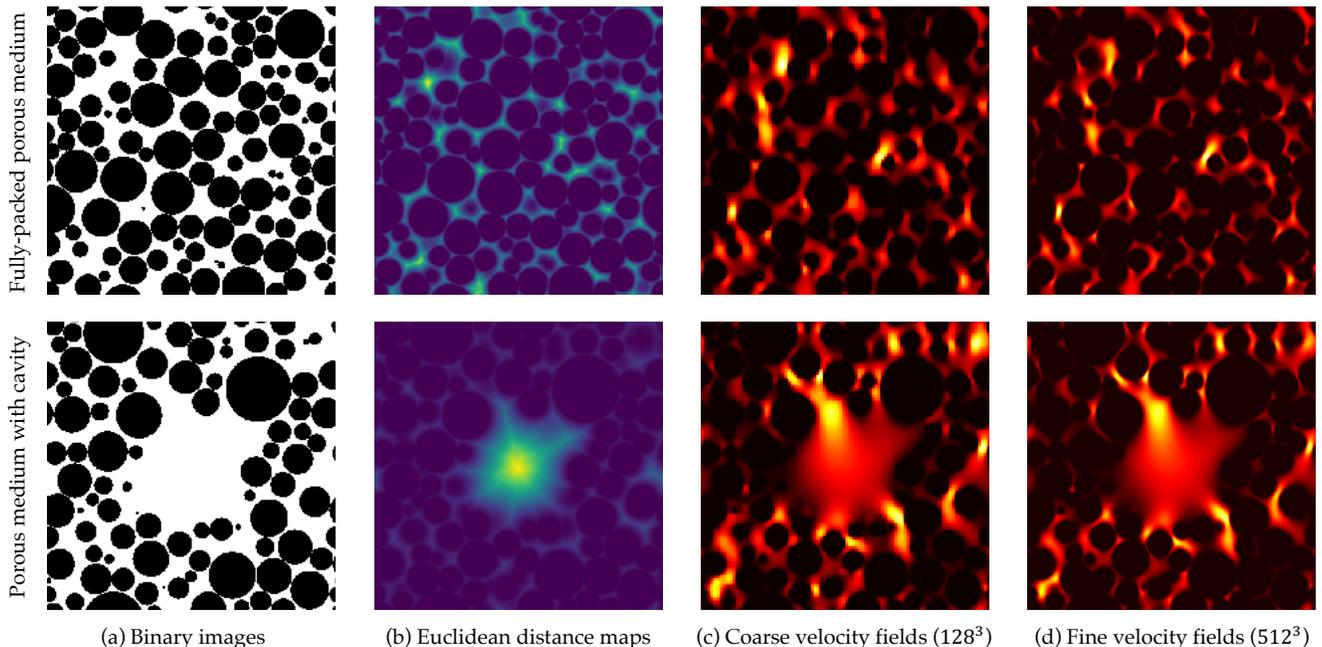


FIG. 1: Cross-sectional view of two different types of porous media and their corresponding geometry representations and velocity fields: (a) binary images, showing pore structures, (b) Euclidean distance maps, showing distance from a point in the pore space to the nearest solid surface, (c) coarse velocity fields, providing coarse global information, and (d) fine-scale velocity fields, used as the model output to train the neural networks. A fully-packed porous medium is shown in the top row whereas a porous medium with a large interior vuggy pore space is shown in the bottom row. Note that all the porous media in this paper are 3D. Here we only show the 2D cross sections for clarity.

160 In this way, we generate the input data (i.e., Euclidean distance maps and coarse velocities) and output data
 (fine-scale velocities) for five training and two testing porous media. However, such large data over the whole domain
 cannot be directly fed into the neural network considering the available GPU memory is limited. Even for a training
 batch size of a single porous medium, it requires about three gigabytes (GB) of memory to locally store the input
 and output, let alone the parameters of the neural network and the calculated gradients. Here, we draw inspiration
 165 from the data processing method in the work of PoreFlow-Net [22] to subsample them into small subvolumes. Given
 that the number of input features in our work is half that for the PoreFlow-Net, the dimension of our subvolumes is
 significantly larger. We evenly segment each porous medium into 64 subvolumes with a dimension of 128^3 (on fine
 mesh), which is shown in Fig. 2. Finally, we have 320 subvolumes for training and 128 for testing.

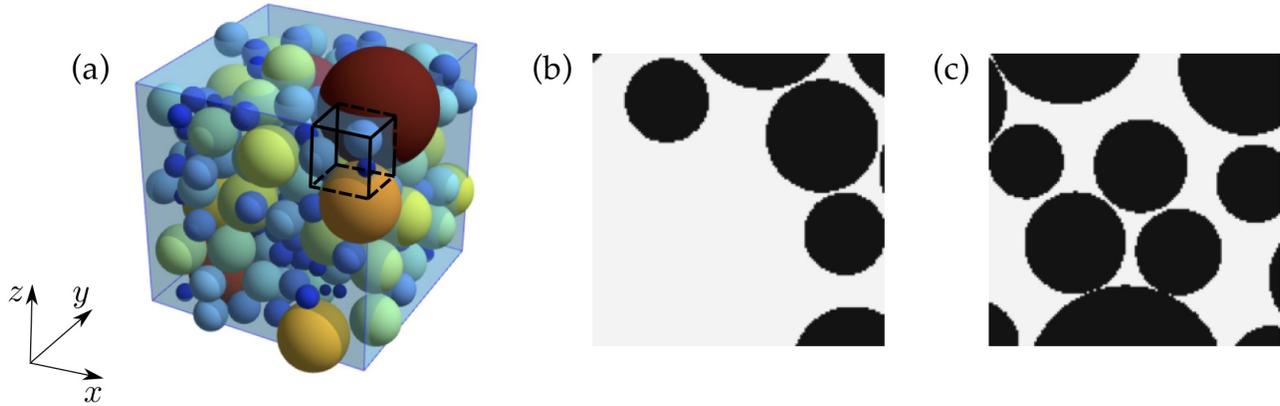


FIG. 2: Method of segmenting a porous medium into smaller subvolumes: (a) large porous medium, packed with spheres, (b) cross section of a subvolume with vuggy pore space, and (c) cross section of a fully-packed subvolume. The large porous medium (512^3) is evenly segmented into 64 subvolumes (128^3). The black cube shown in (a) illustrates the size of a subvolume.

B. U-net with super resolution

170 We propose a CNN-based neural network to predict the fine-scale velocity field in a porous medium based on: (1) fine-scale geometry of pore structure, and (2) coarse velocity field, which is illustrated in Fig. 3.

Similar to the commonly used U-net model [48], the proposed network consists of a contracting path (left) and an expansive path (right), each with four resolution steps. In the contracting path, the 3D image resolution is reduced whereas the information channel number is enhanced, which enables extraction of the pore geometry features. Specifically, each convolutional layer consists of two $3 \times 3 \times 3$ convolutional filters and doubles the number of feature channels followed by a average-pooling layer with a $2 \times 2 \times 2$ filter that coarsens the image resolution by a factor of two in each dimension. In the expansive path, each upsampling layer of the ‘nearest’ mode doubles the resolution of feature maps in each single dimension followed by a concatenation with the corresponding feature maps of the same resolution from the left, contracting path; each convolutional layer consists of two $3 \times 3 \times 3$ convolutional filters and reduces the the number of feature channels. The concatenations aim to reuse the extracted geometry features in the encoder by concatenating them to the levels of the same resolution in the decoder, allowing for the retention of geometric information. In the top level, a $1 \times 1 \times 1$ convolutional filter is finally used to reduce the number of channels to the number of output channels which equals to one in our case. In the neural network, we select the rectified linear activation function (ReLU) as the activation function because it is easy to train and often achieves good performance. In addition, we use batch normalization (BN) [49] before each ReLU operation, which accelerates the training process by reducing the internal covariate shift.

To inform the neural network of global physics, we concatenate the 3D coarse velocity field to the third level in expansive path, where the velocity field is reconstructed level by level. Specifically, the coarse velocity field is fed into a $1 \times 1 \times 1$ convolutional layer to increase the number of feature maps from 1 to 128, which equals to the number of the feature channels of the original maps in the expansive path. In this way, we assume that the geometric information and coarse velocities are of equal importance for the prediction of the fine-scale velocity field.

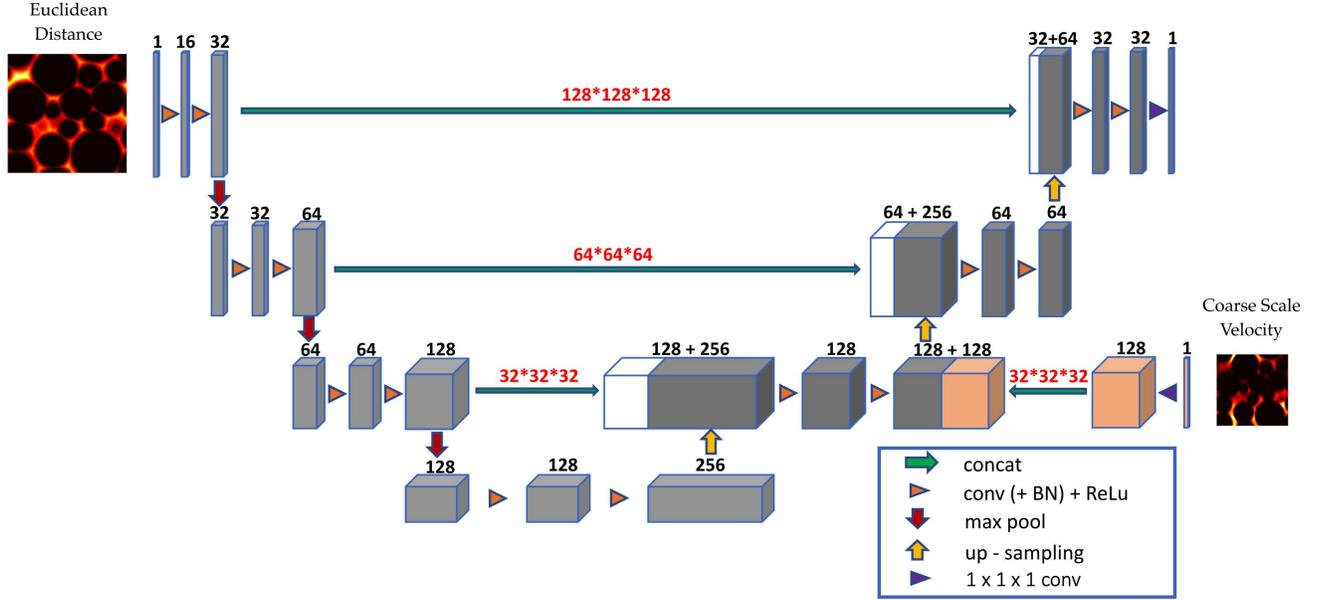


FIG. 3: Schematic plot of 3D U-net assisted with the super resolution technique, which predicts fine-scale velocity fields (128^3) in porous media based on geometries (128^3) and coarse velocity fields (32^3). The gray boxes represent the feature maps and the white boxes represent the copied features mapped from the left side (i.e., the encoder part) of the network. The orange boxes represent the feature maps from the coarse velocity field. The number of feature channels is denoted above each feature map. The resolution of 3D image on each level is denoted above the skip-concatenation arrow. Arrows with different colors indicate different operations, which are defined in the bottom-right figure legend.

Finally, we use the generated 320 input-output data pairs to train the proposed neural network. The input consists of Euclidean distance map (128^3) and coarse velocity field (32^3) and the output is the corresponding fine-scale velocity field (128^3).

III. RESULTS

195

We train the neural networks using two distinct inputs: (1) geometric information only and (2) geometric information and coarse velocity field. Both models are then evaluated on new pore geometries. The prediction accuracy is significantly improved by incorporating the coarse velocity field in the input, especially for those with a large vuggy pore space inside. The neural network is also trained using coarse velocities with an even lower resolution of 16^3 . We can still obtain a much superior prediction result compared to using geometric information alone. In addition, we test the trained model on a real Bentheimer sandstone, whose pore geometry is significantly different from those in the training dataset. The predicted velocity pattern shows high similarity to the ground truth simulated by LBM. In this section, we first show the improvement of predicted fine-scale velocity fields in the porous media by incorporating the coarse velocities in the input (Section A). Then we illustrate the flexibility of the methodology in training the network using coarse velocities with different resolutions (Section B). We further demonstrate the feasibility of the

205

methodology for real rocks by testing on a Bentheimer sandstone (Section C).

A. Neural-network-based prediction aided by coarse velocity field

We first compare the neural-network-based predictions with and without coarse velocity field in the input. The results show that by including coarse velocities in the input, the prediction performance for both training and testing samples is significantly improved. Specifically, different input features are used to train the neural networks: (1) Euclidean distance map only and (2) both Euclidean distance map and coarse velocity field. The training processes are performed on NVIDIA TESLA V100 GPUs using the open source machine learning framework PyTorch [50]. We select the mean squared error (MSE) as the loss function to be optimized, which is defined as:

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N \left\| \mathcal{U}_i - \hat{\mathcal{U}}_i \right\|_2^2, \quad (4)$$

where N is the number of training or testing subvolumes, $\|\cdot\|_2$ denotes the ℓ^2 -norm, $\hat{\mathcal{U}}$ and \mathcal{U} denote the predicted velocity field provided by the neural-network-based model and velocity field simulated using LBM, respectively. The Adam optimizer [51] is adopted to train the neural networks. The training is scheduled such that the learning rate is initialized with 0.001 and is reduced by multiplying a factor of 0.7 every 200 epochs. The entire training process takes 600 epochs with a batch size of four. The training and testing losses for two different scenarios are compared, which is shown in Fig. 4. It is clear that using coarse velocity fields improves the prediction performance, particularly for testing samples. Specifically, when we train the neural network using geometry only, the testing loss converges after 200 epochs and remains nearly constant about 2.3×10^{-4} , which is markedly larger than the final training loss of 7.6×10^{-6} . In contrast, by incorporating the coarse velocities in the input, the testing loss keeps declining to 1.6×10^{-5} , getting much closer to the final training loss of 4.4×10^{-6} .

After training, we evaluate the prediction performance of two trained neural networks on the testing samples. Similarly, the predicted fine velocity fields based on the input with coarse velocities are more accurate than those using geometry only, particularly for the subvolumes with large interior vuggy pore space. We use the prediction error to measure the deviation of predictions from the ground truths, which is defined as the normalized voxel-wise discrepancy between the predicted velocity magnitude \hat{u} and the corresponding ground truth u^* :

$$\text{error} = \frac{\sqrt{\sum_{i=1}^n |\hat{u}_i - u_i^*|^2}}{\sqrt{\sum_{i=1}^n |u_i^*|^2}}, \quad (5)$$

where the summation is performed over all of the n testing voxels (e.g., 128^3 voxels in one subvolume). The prediction errors for all 128 testing samples are calculated to be 9.6% (using coarse velocities) and 38.7% (geometric information only), showing significant improvements aided by the coarse velocity fields. Given that comparing predicted velocities in 3D images is not a good option, we select two samples from all testing subvolumes and show the velocity fields on their cross sections. These two are selected to be representative such that one is fully-packed while the other has the large vuggy pore space. Fig. 5 illustrates a detailed comparison of the predicted velocity fields and their corresponding ground truths on the cross sections of both subvolumes. We can see that the predictions based on Euclidean distance maps only diverge significantly from the ground truths, especially for the subvolume with vuggy pore space, in which

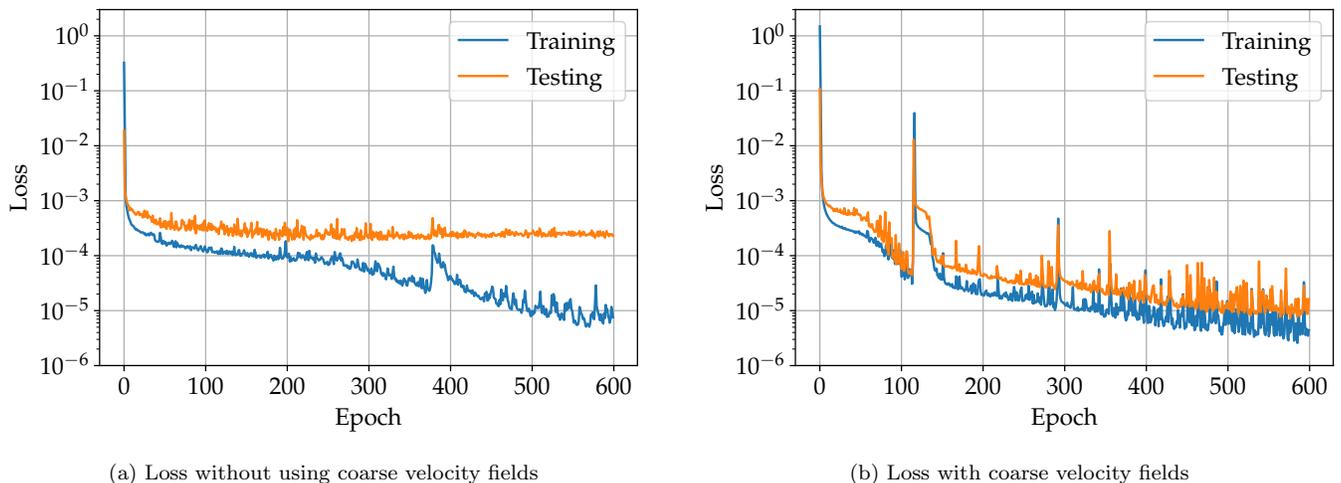


FIG. 4: Training and testing losses for two different scenarios: (a) training using Euclidean distance maps only, and (b) training using both Euclidean distance maps and coarse velocity fields. By using geometric information alone, the testing loss shows rapid convergence and then remains at a value noticeably larger than the training loss. After incorporating the coarse velocity fields in the input, the testing loss continues declining and gets close to the training loss.

the predicted velocity field exhibits a very different pattern with non-physical velocity discontinuity and a much smaller velocity magnitude. When coarse velocities are included in the input, the predictions for both subvolumes are improved substantially and become pretty comparable to the ground truths. Specifically, the prediction error decreases from 33.2% to 8.9% for the fully-packed subvolume (top row) and more dramatically from 51.6% to 7.5% for that with vuggy pore space (bottom row).

From the observation above, the vuggy pore space seems to have a significant impact on the prediction performance of the model which is trained using geometric information only. We further investigate the impact by altering the volume of the vuggy pore space therein. The results demonstrate that as the vuggy volume grows, predicting fine-scale velocities based on geometry only becomes more challenging. In contrast, including coarse velocities in the input significantly improves the prediction accuracy, especially for the predicted velocity field in vuggy pore space. We change the radius of the spherical vuggy pore space in a fully-packed porous medium ($\phi_{\text{target}} = 0.33$) to generate four testing porous media with different vuggy volumes, as illustrated in Fig. 6(a). The values of four different radii are 0.15, 0.19, 0.22 and 0.24, with the corresponding vuggy volume percentages of 1.4%, 2.9%, 4.5% and 5.8%. The vuggy volume percentage is defined as the volume of the spherical vuggy pore space divided by the volume of the entire domain. We evaluate the trained neural networks on these porous media and calculate both global and local prediction errors. Note that for each porous medium, we first segment it into 64 small subvolumes for testing; predicted velocities in each subvolume are then pieced together as the prediction in the entire domain. The global prediction error is calculated over the entire domain, while the local is calculated over a smaller region which is dominated by the large vuggy pore space. The local region has half the length of the domain and is denoted by a dashed square in Fig. 6(a). The prediction errors for different vuggy volumes are shown in Fig. 6(b). We can see that when the coarse velocities are not used for prediction, both global and local errors increase as the vuggy volume grows, implying that

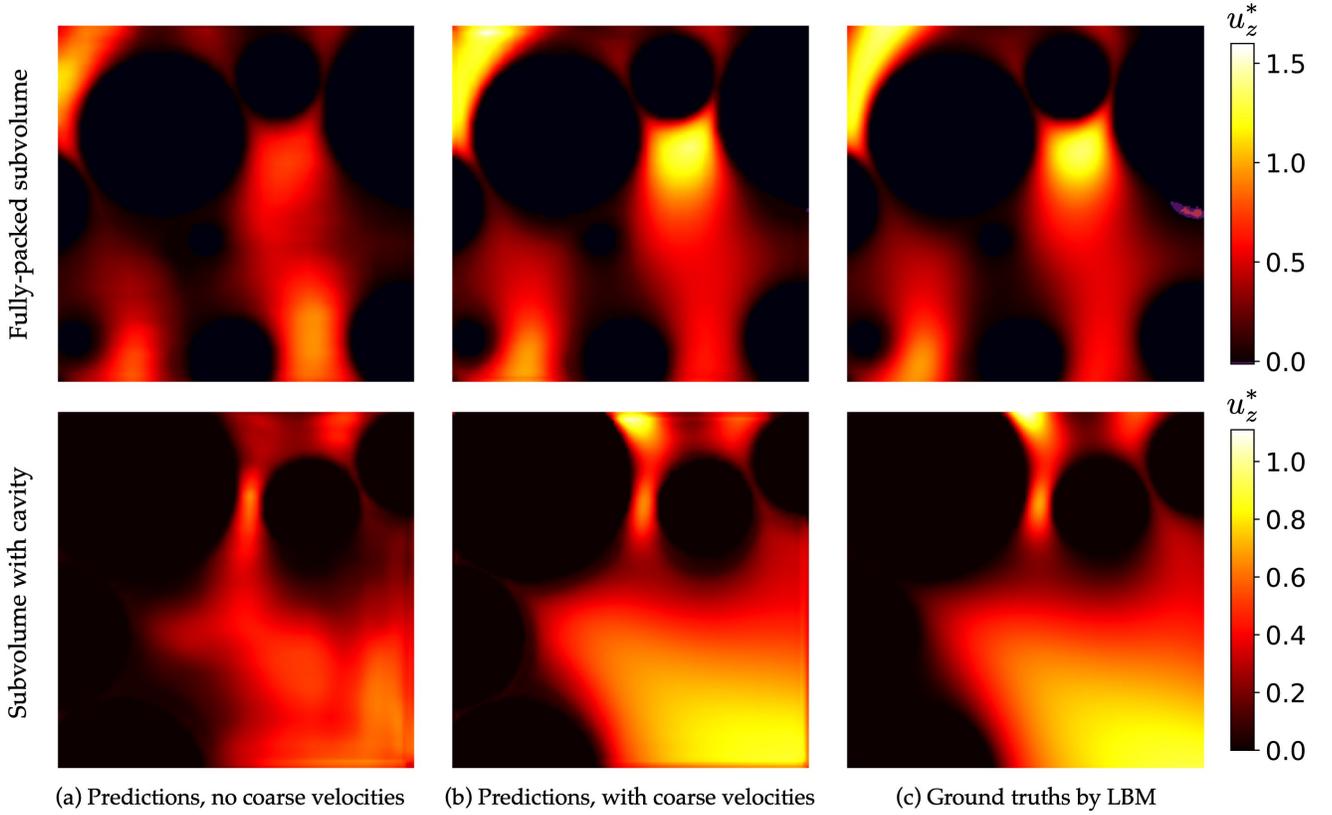
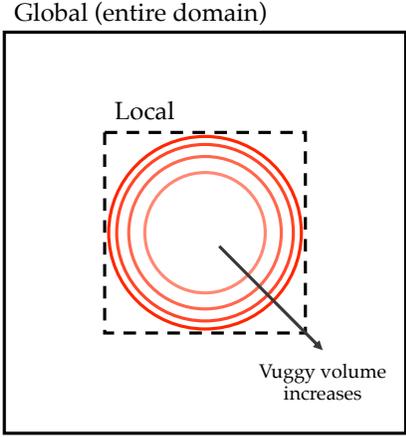


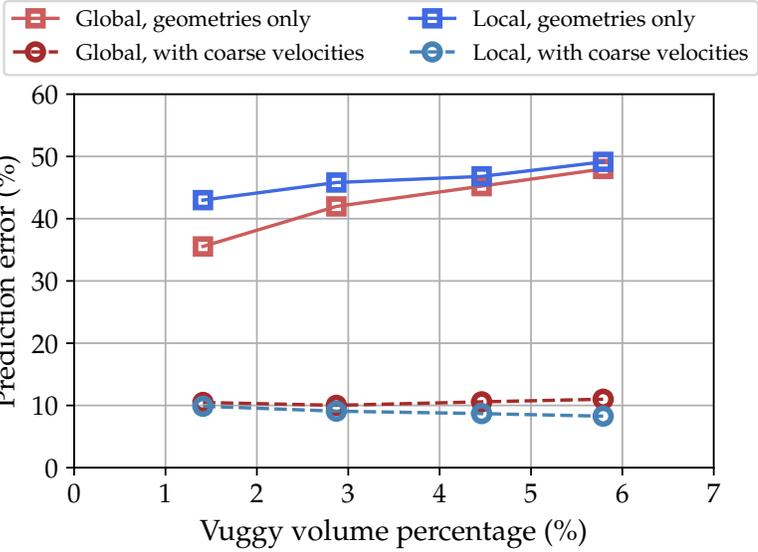
FIG. 5: Comparison of the predicted fine-scale velocity fields provided by different trained models with the corresponding ground truths in two testing subvolumes: (a) predictions based on geometric information only, (b) predictions using both geometric information and coarse velocity fields, and (c) ground truths simulated using the LBM. The top row is for the fully-packed testing subvolume whereas the bottom row is for the testing subvolume having an interior vuggy pore space. Note that here we show the 2D cross sections of velocity fields for clarity.

a larger vuggy pore space complicates the prediction. Additionally, the local error is clearly larger than the global error, indicating that the interior large vuggy pore space is the primary source of the inaccuracy. After incorporating
 260 the coarse velocities in the input, both global and local errors for four vuggy volumes reduce dramatically and reach a value about 10%. This improvement in prediction accuracy indicates that the trained model is capable of accurately predicting the fine-scale velocity fields in porous media with changing vuggy pore space and is not sensitive to vuggy volume. What is remarkable is that when the vuggy volume grows larger, the local error reduces. This can be interpreted that the predicted velocity field in large vuggy pore space is much easier to improve with the help of
 265 coarse velocities as compared to the flow through densely packed spheres.

In the discussion above, the trained neural networks are tested purely on subvolumes which lack global geometric information. We further eliminate the ill-posedness by using two full (i.e., non-subsampled) porous assemblies and then investigate the prediction performance of the trained models on these two assemblies. The results show that even with complete geometric information of the whole domain, the geometry-based model still cannot provide a
 270 reasonable prediction. These two assemblies are created such that one is fully-packed with spheres and the other has a large vuggy pore space in the center. We keep the size of both assemblies the same as the testing subvolumes, but



(a) Porous media with different vuggy volumes



(b) Prediction errors for different vuggy volumes

FIG. 6: Cross sections of porous media with increasing vuggy volumes and the corresponding prediction errors. (a) Porous media with four different vuggy volumes. Circles, colored from light red to dark red, denote the increasing regions of the interior spherical vuggy pore space. The black solid and dashed squares represent the global (i.e., entire) and local domains. (b) Global and local prediction errors against increasing vuggy volume percentage. Two models are compared: one is trained with geometric information only whereas the other is trained with both geometric information and coarse velocity fields.

significantly increase the number and change the distribution of the spheres contained. The predicted velocity fields based on global geometries alone are still largely different from the ground truths, with prediction errors of 57.2% for the fully-packed assembly and 43.4% for the other with a interior vuggy pore space. This is reasonable because extrapolating the learned highly nonlinear functions to new porous media with different pore structures can be a very challenging task. Such weak performance can be improved by incorporating coarse velocities in the input, which is shown in Fig. 7. Specifically, the overestimations of velocity magnitude along the bottom and right borders in the fully-packed assembly (top row), as well as the deviations in the assembly with the large vuggy pore space (bottom row), are both reduced. Accordingly, the prediction errors drop significantly to 22.5% and 17.7%.

Additionally, we conduct a parametric study and find that the number of training samples (320) used in this work is sufficient and necessary to learn the neural network. The number of training samples is first halved to be 160 by taking one from every two current samples. Such trained neural network is then evaluated on the original 128 testing subvolumes. The prediction error for all testing samples increases to 13.4%, which is larger than 9.6% when using 320 training samples. We also increase the number of training samples by a factor of two to have 640 training samples. The additional 320 samples are still taken from the original five training porous media but obtained by random segmentations. The model trained with 640 samples is then evaluated on the same 128 testing subvolumes and the prediction error is 9.2%, showing no significant improvement compared to the model trained with 320 samples.

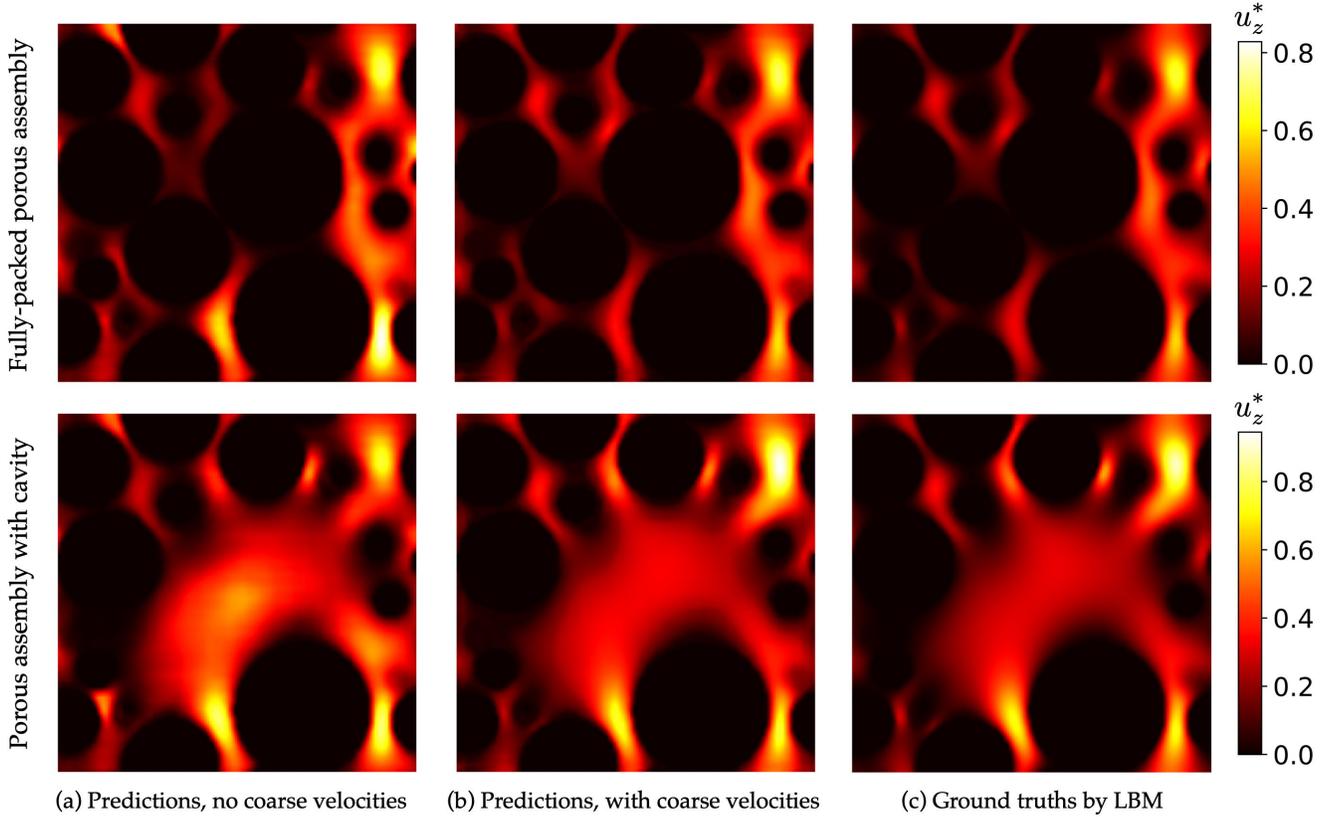


FIG. 7: Comparison of the predicted fine-scale velocity fields provided by different trained models with the corresponding ground truths in another two porous assemblies: (a) predictions based on geometric information only, (b) predictions using both geometric information and coarse velocity fields, and (c) ground truths simulated by LBM. Note that the utilized porous assemblies are both full porous media, meaning that they are not obtained by segmentation. The top row is for the fully-packed porous assembly and the bottom row for that with interior vuggy pore space.

Therefore, we conclude that 320 training samples are sufficient to learn the neural-network-based model, at least for the testing data generated in this work.

B. Flexibility in using coarse velocity fields with various resolutions

Another advantage of the proposed method is the flexibility of the coarse-velocity resolution. In other words, the neural network can be trained using coarse velocity fields with various resolutions. Even with a very coarse velocity field, the predictions still surpass those based on geometric information only.

In the discussion above, we coarsen the fine mesh (128^3) by a factor of four and simulate the velocity field on the coarse mesh (32^3), which serves as the coarse velocity field to the network. The coarse-velocity resolution, however, can be varied thanks to the multi-level architecture of the network. One method is to concatenate coarse velocities with varying resolutions to the corresponding levels of the network. While the operation is straightforward, it will inevitably change the architecture of the neural network. For example, coarse velocities with a resolution of 64^3 must

be fed into the second level of the decoder, while those with a resolution of 16^3 must be sent into the bottom level. Another alternative is to maintain the current architecture but downsample or upsample different coarse-velocity resolutions to 32^3 . This strategy may result in the information loss but it enables easy refinement of trained networks with varied coarse-velocity resolutions.

Here we keep the network unchanged (i.e., same as shown in Fig. 3), but use the coarse velocity fields with an even lower resolution of 16^3 . To this end, the coarse velocities are preprocessed: the resolution of the coarse velocities is doubled in each direction via the ‘nearest’ upsampling mode, which pads the neighboring eight (i.e., 2^3) voxels with the same value. Given that we have already obtained the model trained with coarse velocities (32^3), the upsampled coarse velocities, together with their geometries, can be directly fed into this model to predict the fine-scale velocities. The predictions of the same two testing subvolumes are shown in Fig. 8(b). We can observe that even using the coarse velocities with an lower resolution, the prediction patterns improve significantly as compared with the baseline predictions in Fig. 8(a). However, there are clear layer artifacts (i.e., unsmoothness) in the predicted velocity fields, which may be caused by the preprocessing operation to increase the coarse-velocity resolution from 16^3 to 32^3 .

To remove the unsmoothness, we retrain the neural network using the coarse velocities with a resolution of 16^3 . Specifically, the original 320 subvolumes still serve as training samples; the neural network is trained using Euclidean distance maps and upsampled coarse velocity fields. We maintain all the training settings unchanged and evaluate the retrained neural network on the same two testing subvolumes. The prediction results provided by the retrained model are shown in Fig. 8(c). It is evident that the unsmoothness is eliminated when the coarse velocities with a resolution of 16^3 are used for retraining. Accordingly, the prediction errors of these two subvolumes are lowered from 31.1% to 14.9% for the fully-packed one and from 22.2% to 15.3% for the other with vuggy pore space.

The prediction accuracy for these two testing subvolumes is more clearly visualized in Fig. 9, which shows the xy -plane-averaged velocity profiles along z direction. For both subvolumes, we can see that the predictions (grey) using geometry only differ significantly from the ground truth (black). This is particularly obvious for the subvolume with large vuggy pore space, in which the trained model is incapable of predicting the flow through the cavity space (i.e. $z/\ell_c \in [0.6, 1.0]$). After training the neural networks with both geometry and coarse velocities as inputs, the predicted velocity profiles (green for 32^3 and red for 16^3 , both solid lines) approach the ground truths markedly. It is clear that the coarse velocities with a resolution of 16^3 improve the predicted velocity profile more than those with a higher resolution of 32^3 . This appears to contradict the fact that the higher-resolution velocities retain more physics of the flow. However, it is conceivable because the averaging operation over the xy planes may obscure the discrepancy between the predicted profiles based on the coarse velocities (16^3) and the ground truths. The brown curve exhibits clear overestimation. This is acceptable because we train the network with higher-resolution coarse velocities but test it on new subvolumes using lower-resolution coarse velocities. In addition, we compare the neural-network-based predictions with the reconstructed fine-scale velocities obtained by tricubic interpolation from coarse velocities. The light green and red dashdot lines represent the interpolated velocity fields based on coarse velocities with resolutions of 32^3 and 16^3 , respectively. As can be seen, the tricubic interpolation is incapable of producing a reasonable velocity field in terms of the velocity magnitude. By comparison, the CNN-based super resolution, when combined with the fine-scale pore structure images, significantly outperforms the simple tricubic interpolation to a finer mesh. More details of the velocity fields simulated on coarse meshes and the corresponding interpolated fine-scale velocity fields can be found in Appendix A. Furthermore, the prediction errors for testing samples based on different input features

and trained models are presented in Table I.

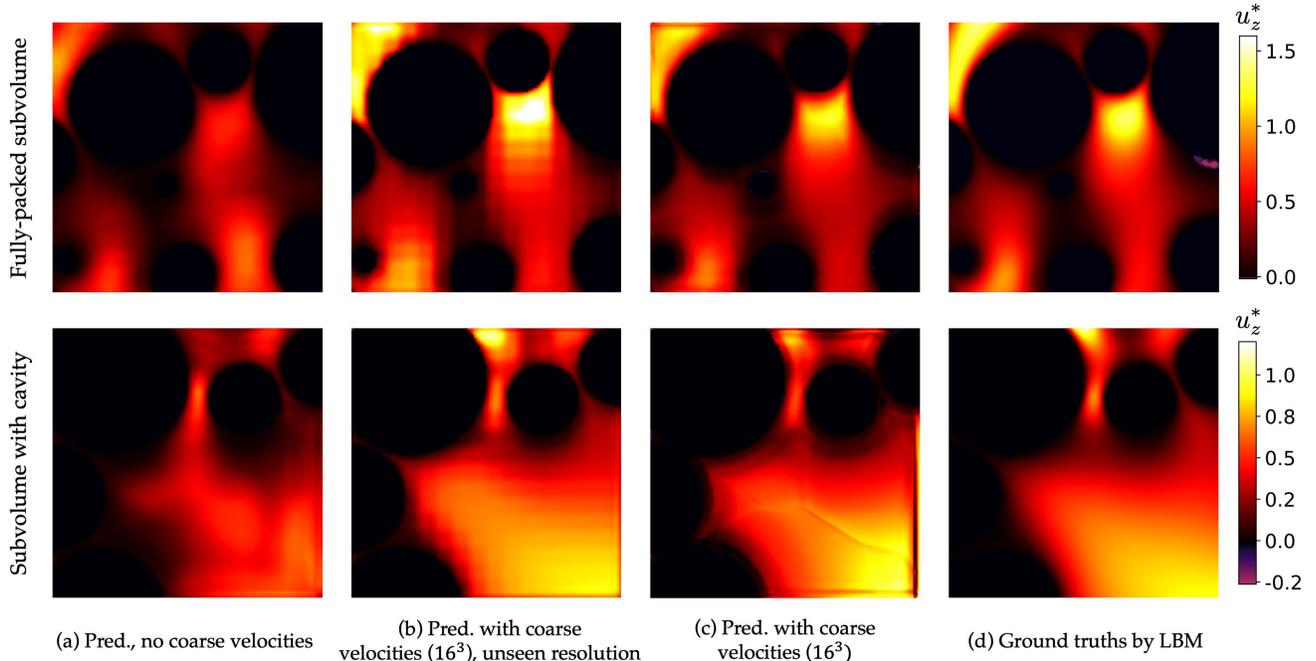


FIG. 8: Comparison of the predicted fine-scale velocities based on different trained models and inputs with the corresponding ground truths: (a) predictions using geometric information only, (b) predictions using coarse velocities with a lower resolution (16^3), but provided by the model trained with coarse velocities with a higher resolution of 32^3 , (c) predictions using coarse velocities with a lower resolution (16^3), provided by the model retrained with coarse velocities with the same resolution of 16^3 , and (d) corresponding ground truths by LBM. Note that (b) and (c) use the same input features but different trained models. The top row is for the fully-packed subvolume and the bottom is for the subvolume with interior large vuggy pore space.

C. Prediction on real Bentheimer Sandstone

340 We further test the previously obtained model, which is trained using coarse velocities (32^3), on a real Bentheimer sandstone to demonstrate the feasibility of our method. We choose the Bentheimer sandstone because it is well resolved and often used as a standard benchmark rock [52, 53]. To be more precise, we select a Bentheimer sample with a dimension of $1600 \times 900 \times 900$ from the Digital Rocks Portal [54], and use a subsection of the complete image to represent the testing rock with a dimension of 512^3 , as shown in Fig. 10. We process the binary image to a Euclidean

345 distance map and simulate both coarse- and fine-scale velocity fields therein using the LBM. The corresponding parameters in MRT are specified identically to those in Section II A. Likewise, we evenly segment the geometry of the rock, as well as the velocity fields, into 64 small subvolumes for testing. The distance maps and coarse velocities of all subvolumes are loaded in one batch and fed into the well-trained model to predict the fine-scale velocity fields.

350 With the coarse velocity fields, the predicted fine-scale velocities are comparable to the simulation results using LBM in terms of flow patterns. Three subvolumes are selected randomly for a detailed comparison. For each subvolume,

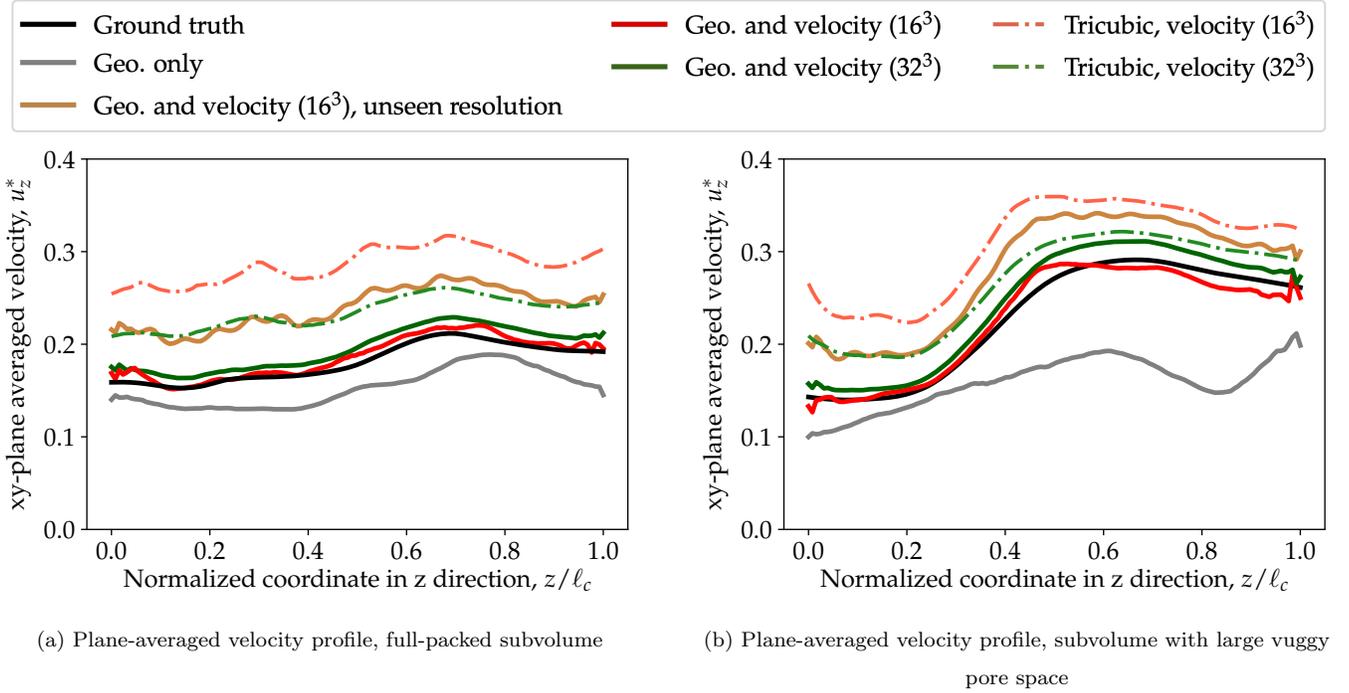


FIG. 9: Comparison of the xy -plane-averaged velocity profiles along z axis based on different inputs and trained models with the corresponding ground truths in two testing subvolumes: (a) velocity profiles in a fully-packed subvolume and (b) velocity profiles in a subvolume with large vuggy pore space. The z coordinate is normalized by the length of the subvolume ℓ_c such that it ranges from 0 to 1. Note that “unseen resolution” refers to the case where the coarse velocities with a low resolution of 16^3 are used in the input for the model that is trained using the coarse velocities with a higher resolution of 32^3 . The dashdot lines represent the profiles of the fine-scale velocity fields which are obtained via tricubic interpolation from various coarse-scale velocity fields.

a xy -cross section is chosen such that the pore space on this cross section is big enough and thus the velocity fields can be clearly visualized. The predicted fine-scale velocity fields on these three cross sections are compared to the baseline predictions using geometries only and the corresponding ground truths, as illustrated in Fig. 11. When coarse velocities are included in the input for training, we can see the predictions get significant improvement by comparison to the baseline velocity fields. The flow patterns inside these three subvolumes are quite similar to the ground truths, despite the trained model has never seen such geometries with non-spherical grains. However, there are still differences between the predictions and ground truths in terms of velocity magnitude. This is unsurprising because the pore structures of Bentheimer sandstone is significantly more complex than those of training subvolumes packed with spherical grains. The prediction error of 25.1% for all 64 testing subvolumes also explains such difficulty of extrapolation. Nonetheless, it is expected that increasing the amount of real data for training the neural network will improve its prediction capability on real rocks.

TABLE I: Prediction errors for testing subvolumes using different input features. The first four columns provide the prediction errors of neural-network-based results, while the last two show the prediction errors of the tricubic interpolated velocities. The fully-packed subvolume and the subvolume with large vuggy pore space are selected from 128 testing samples and used for flow visualizations throughout this paper.

Prediction error (%)	Geo.	Vel. (32^3) & Geo.	Vel. (16^3) & Geo.	Vel. (16^3) & Geo., unseen resolution	Vel. (32^3), tricubic	Vel. (16^3), tricubic
Fully-packed subvolume	33.2	8.9	14.9	31.1	31.4	58.0
Subvolume with cavity	51.6	7.5	15.3	22.2	19.1	36.4
All 128 testing samples	38.7	9.6	16.0	32.5	33.2	62.1

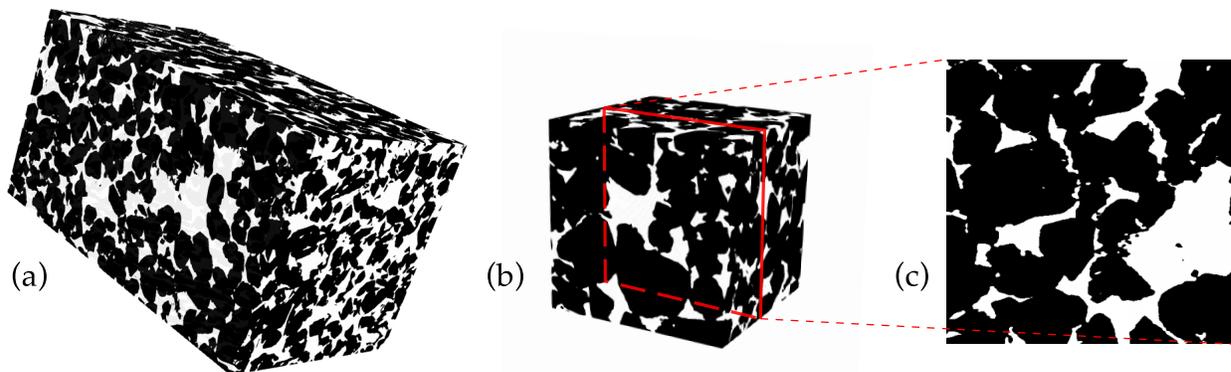


FIG. 10: 3D binary image of a real Bentheimer sandstone used as the testing sample: (a) 3D binary image of the entire Bentheimer sandstone with dimensions of $1600 \times 900 \times 900$ voxels, (b) a subsample with dimensions of 512^3 voxels, which is used as the final testing rock, and (c) an enlarged 2D cross section of the testing rock, showing the pore structure in detail.

IV. CONCLUSIONS

Direct numerical simulations are commonly used to simulate the pore-scale velocity fields in porous media. However, such method is expensive and time-consuming. Drawing inspirations from the super resolution technique, we incorporate the coarse velocity field in the input and predict the fine-scale velocities through porous media using both the geometry of pore structures and the coarse velocities. The coarse velocity field is obtained by solving the Stokes equation on a coarsened mesh that it still retains significant physics of the fine-scale flow to aid in prediction. Additionally, it informs the network of global physics and thus helps to regularize the ill-posedness of the learning problem due to the use of local geometry. The proposed U-net-based neural networks are trained with sphere-packed porous media and tested on different porous media including even real rocks. The results demonstrate the significant improvement of prediction by having the coarse velocity field in the input, as well as the training flexibility of using

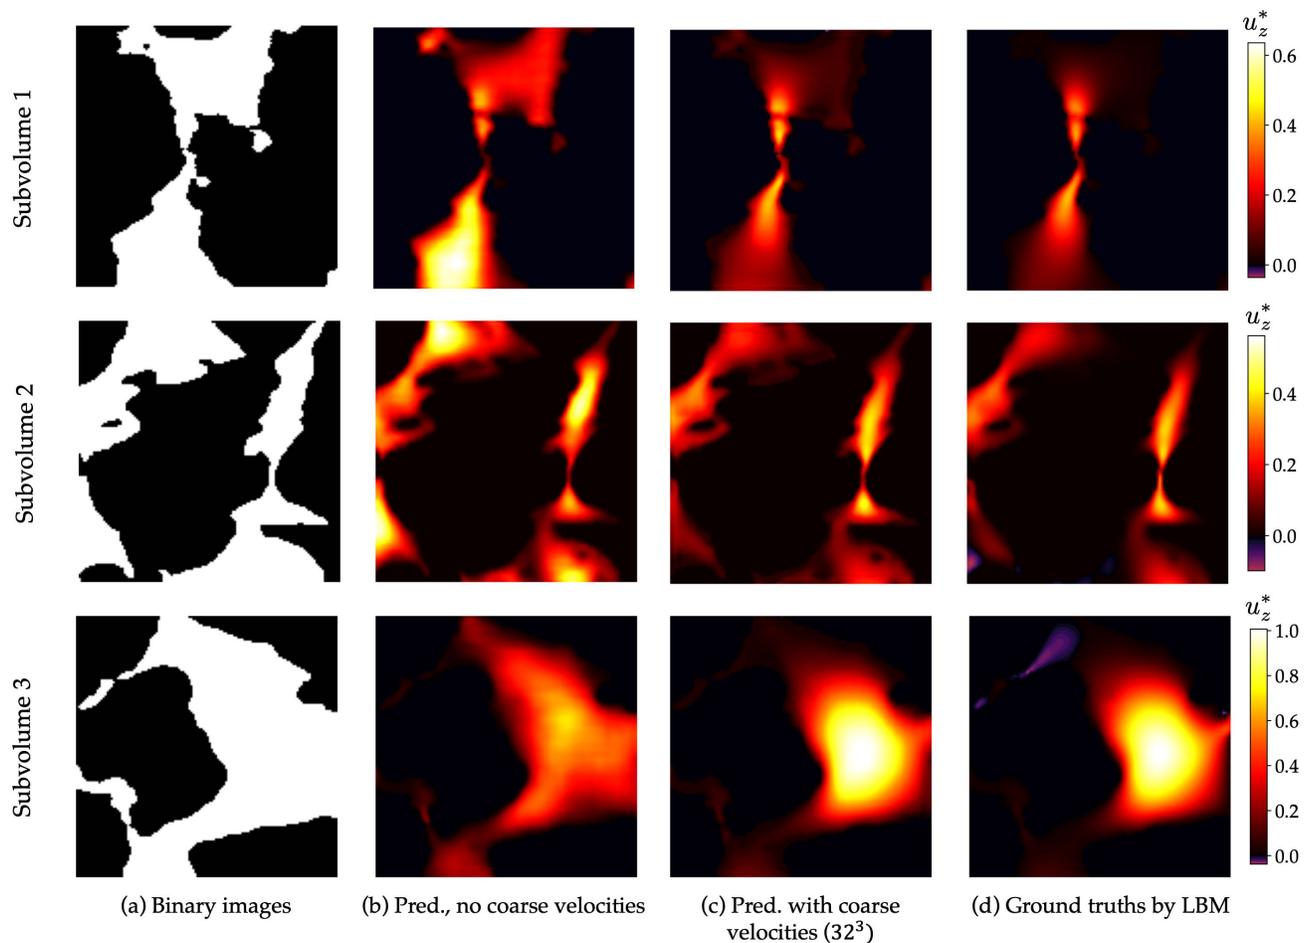


FIG. 11: Comparison of the predicted fine-scale velocity fields (128^3) with the corresponding ground truths in three real Bentheimer sandstone subvolumes: (a) binary images of the rocks, (b) predictions based on geometric information only, (c) predictions using both geometric information and coarse velocities (32^3), and (d) ground truths by LBM. The velocity fields are illustrated on 2D cross sections for clear visualization.

coarse velocities with varying resolutions and the extrapolation capability of the trained model for real rocks.

Despite the preliminary successes demonstrated in this work, there are still several directions that can be explored. First and foremost, this paper discusses only single-phase flow through porous media. However, multiphase flow is more prevalent in real-world situations [5, 55] and warrants further investigation. Our method should still be applicable to predicting fine-scale multiphase velocity fields in porous media. This is because multiphase flow is significantly more complicated than single-phase flow, necessitating a greater reliance on physics-based information for effective prediction. This information can be embedded inexpensively by incorporating the coarse velocities in the input. Moreover, we train the proposed neural network solely based on the labeled data, without imposing any physical constraint. Future research could look into ways to impose physical laws (e.g., mass conservation) for the prediction. The work of embedding hard physical constraints for 3D turbulence provides a possible approach to achieve this goal by using non-trainable layers with physics [56]. Finally, the core concept of this study may be reliably extended to related issues by including coarse information in the input. Both the input and output are in a very high-dimensional

space in these issues, and the output is highly sensitive to changes in the input, rendering the trained model unsuitable
 385 for extrapolation tasks. Informing neural networks of coarse information that is collected inexpensively but contains
 critical information about the output can be an effective strategy for increasing the model’s generalization capability.

DATA AVAILABILITY

The code for data generation and neural network training will be made available in a public Github repository [57],
 and thus the results in this paper can be straightforwardly reproduced and further developed by the readers.

ACKNOWLEDGMENTS

The authors appreciate the editor and reviewers’ feedback and suggestions, which helped in improving the quality
 of the article. This work was financially supported by the University Coalition for Fossil Energy Research (UCFER)
 Program under the U.S. Department of Energy’s National Energy Technology Laboratory through the Award Number
 DE-FE0026825 and Subaward Number S000038-USDOE. The computational resources used for this project were
 395 provided by the Advanced Research Computing (ARC) of Virginia Tech, which is gratefully acknowledged.

Appendix A: Reconstruction of fine-scale velocity field by tricubic interpolation

In this work, we consider using an interpolated fine-scale velocity field as another baseline prediction to demonstrate
 the superiority of our proposed methodology. This baseline velocity field is reconstructed from a coarse velocity field
 using tricubic interpolation, which is a widely used technique for image upsampling in computer graphics. Tricubic
 400 interpolation is an extension of bicubic interpolation for interpolating data points on a 3D regular grid. Specifically,
 for any unknown data point in the porous medium, tricubic interpolation takes 64 neighboring pixels ($4 \times 4 \times 4$) into
 account, which achieves a smoother image with fewer interpolation artifacts as compared to trilinear interpolation.

In Fig. 12, we show the simulated velocity fields on two coarsened meshes (16^3 and 32^3) and the corresponding
 reconstructed fine-scale velocity fields (128^3) via tricubic interpolation for two testing subvolumes. As can be seen,
 405 the patterns of the interpolated velocity fields closely match the ground truths, owing to the coarse velocities retaining
 critical flow physics. However, the prediction errors remain relatively large. When coarse velocities (16^3) are used for
 interpolation, the prediction errors for the fully-packed subvolume and the subvolume with interior vuggy pore space
 are 58.0% and 36.4%, respectively. When coarse velocities with a higher resolution 32^3 are used, the prediction errors
 decrease to 31.4% and 19.1%. The interpolation-based prediction errors are reasonable because upsampling by simply
 410 considering neighboring pixels is insufficient for reconstructing the globally dependent flow physics. Even for the solid
 grains of pore structure, interpolation alone cannot adequately approximate them at fine scale. It should be noted
 that we embed the original grain pixels into the interpolated images to obtain a more realistic baseline prediction for
 comparison. This method is used to calculate all the above interpolation-based prediction errors.

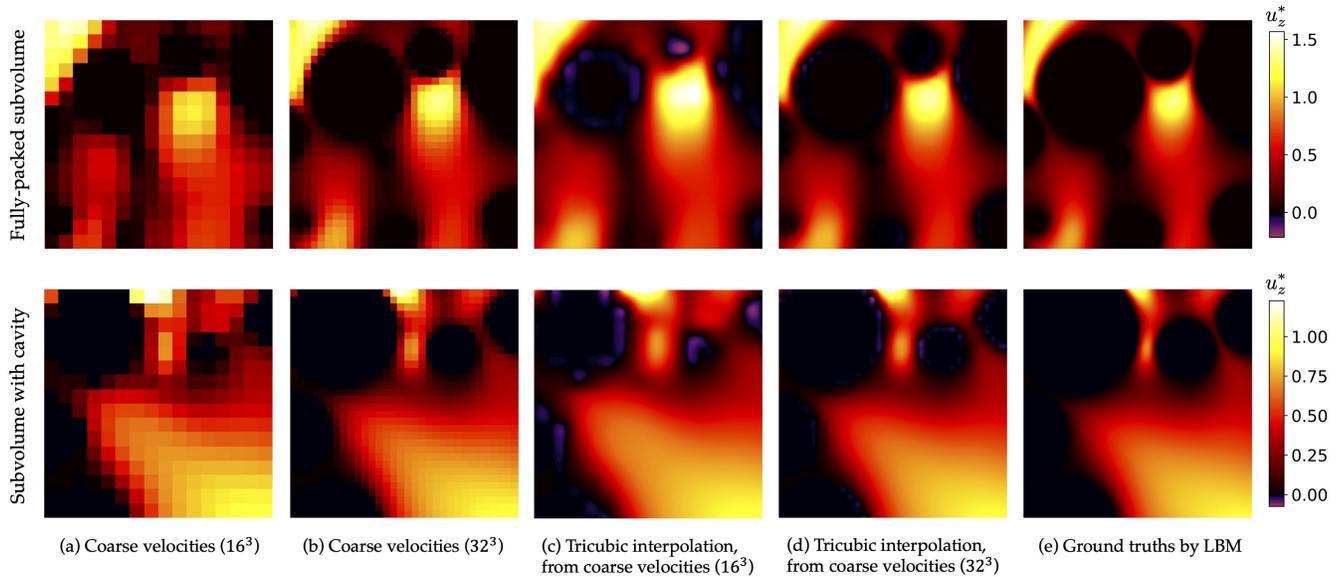


FIG. 12: Cross-sectional view of the coarse velocity fields in two testing subvolumes with different resolutions and the corresponding reconstructed fine-scale velocities via tricubic interpolation: (a) coarse velocities with a resolution of 16^3 , (b) coarse velocities with a resolution of 32^3 , (c) reconstructed velocity fields via tricubic interpolation from the coarse velocities (16^3), (d) reconstructed velocity fields via tricubic interpolation from the coarse velocities (32^3), and (e) fine-scale velocity fields simulated by LBM. The top row is for the fully-packed subvolume and the bottom is for the subvolume with interior large vuggy pore space.

-
- [1] M. J. Blunt, B. Bijeljic, H. Dong, O. Gharbi, S. Iglauer, P. Mostaghimi, A. Paluszny, and C. Pentland, Pore-scale imaging and modelling, *Advances in Water Resources* **51**, 197 (2013).
 415
- [2] J. Bear, *Dynamics of Fluids in Porous Media* (Courier, 2013).
- [3] M. Hilpert and C. T. Miller, Pore-morphology-based simulation of drainage in totally wetting porous media, *Advances in Water Resources* **24**, 243 (2001).
- [4] Q. Kang, P. C. Lichtner, and D. Zhang, An improved lattice Boltzmann model for multicomponent reactive transport in porous media at the pore scale, *Water Resources Research* **43** (2007).
 420
- [5] J. E. McClure, J. F. Prins, and C. T. Miller, A novel heterogeneous algorithm to simulate multiphase flow in porous media on multicore CPU-GPU systems, *Computer Physics Communications* **185**, 1865 (2014).
- [6] M. A. Spaid and F. R. Phelan Jr, Lattice Boltzmann methods for modeling microscale flow in fibrous porous media, *Physics of Fluids* **9**, 2468 (1997).
- [7] E. S. Boek and M. Venturoli, Lattice-Boltzmann studies of fluid flow in porous media with realistic rock geometries, *Computers & Mathematics with Applications* **59**, 2305 (2010).
 425
- [8] P. Jenny, S. Lee, and H. A. Tchelepi, Multi-scale finite-volume method for elliptic problems in subsurface flow simulation, *Journal of Computational Physics* **187**, 47 (2003).
- [9] Y. Song, K. Chung, T. Kang, and J. Youn, Prediction of permeability tensor for three dimensional circular braided preform by applying a finite volume method to a unit cell, *Composites Science and Technology* **64**, 1629 (2004).
 430

- [10] Z. Sun, R. E. Logé, and M. Bernacki, 3D finite element model of semi-solid permeability in an equiaxed granular structure, *Computational Materials Science* **49**, 158 (2010).
- [11] C. Sandino, P. Krociczek, D. D. McErlain, and S. K. Boyd, Predicting the permeability of trabecular bone by micro-computed tomography and finite element modeling, *Journal of Biomechanics* **47**, 3129 (2014).
- 435 [12] K. Duraisamy, G. Iaccarino, and H. Xiao, Turbulence Modeling in the Age of Data, *Annual Review of Fluid Mechanics* **51**, 357 (2019).
- [13] J.-L. Wu, H. Xiao, and E. Paterson, Physics-informed machine learning approach for augmenting turbulence models: A comprehensive framework, *Physical Review Fluids* **3**, 074602 (2018).
- [14] J. Han, L. Zhang, R. Car, and E. Weinan, Deep potential: A general representation of a many-body potential energy
440 surface, *Communications in Computational Physics* **23**, 629 (2018).
- [15] L. Zhang, J. Han, H. Wang, R. Car, and E. Weinan, Deep Potential Molecular Dynamics: A Scalable Model with the Accuracy of Quantum Mechanics, *Physical Review Letters* **120**, 143001 (2018).
- [16] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, *et al.*, Highly accurate protein structure prediction with AlphaFold, *Nature* **596**, 583 (2021).
- 445 [17] N. Srisutthiyakorn, Deep-learning methods for predicting permeability from 2D/3D binary-segmented images, in *SEG Technical Program Expanded Abstracts 2016* (Society of Exploration Geophysicists, 2016) pp. 3042–3046.
- [18] J. Wu, X. Yin, and H. Xiao, Seeing permeability from images: fast prediction with convolutional neural networks, *Science Bulletin* **63**, 1215 (2018).
- [19] J. Kozeny, Über kapillare leitung der wasser in boden, Royal Academy of Science, Vienna, Proc. Class I **136**, 271 (1927).
- 450 [20] P. C. Carman, Fluid flow through granular beds, *Trans. Inst. Chem. Eng.* **15**, 150 (1937).
- [21] Y. D. Wang, T. Chung, R. T. Armstrong, and P. Mostaghimi, ML-LBM: Predicting and Accelerating Steady State Flow Simulation in Porous Media with Convolutional Neural Networks, *Transport in Porous Media* **138**, 49 (2021).
- [22] J. E. Santos, D. Xu, H. Jo, C. J. Landry, M. Prodanović, and M. J. Pyrcz, PoreFlow-Net: A 3D convolutional neural network to predict fluid flow through porous media, *Advances in Water Resources* **138**, 103539 (2020).
- 455 [23] P. Wesseling, *Introduction to multigrid methods*, Tech. Rep. (1995).
- [24] C. Dong, C. C. Loy, K. He, and X. Tang, Image Super-Resolution Using Deep Convolutional Networks, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**, 295 (2015).
- [25] K. Fukami, K. Fukagata, and K. Taira, Super-resolution reconstruction of turbulent flows with machine learning, *Journal of Fluid Mechanics* **870**, 106 (2019).
- 460 [26] B. Liu, J. Tang, H. Huang, and X.-Y. Lu, Deep learning methods for super-resolution reconstruction of turbulent flows, *Physics of Fluids* **32**, 025105 (2020).
- [27] A. Subramaniam, M.-L. Wong, R. Borker, S. Nimmagadda, and S. Lele, Turbulence Enrichment with Physics-informed Generative Adversarial Network, in *Neural Information Processing Systems* (2020).
- [28] Y. D. Wang, R. T. Armstrong, and P. Mostaghimi, Enhancing Resolution of Digital Rock Images with Super Resolution Convolutional Neural Networks, *Journal of Petroleum Science and Engineering* **182**, 106261 (2019).
- 465 [29] J. E. Santos, Y. Yin, H. Jo, W. Pan, Q. Kang, H. S. Viswanathan, M. Prodanović, M. J. Pyrcz, and N. Lubbers, Computationally Efficient Multiscale Neural Networks Applied to Fluid Flow in Complex 3D Porous Media, *Transport in Porous Media* **140**, 241 (2021).
- [30] J. E. McClure, SpherePackTools, <https://github.com/JamesEMcClure/SpherePackTools> (2015).
- 470 [31] Scipy, <https://scipy.org/> (2001).
- [32] S. Kim and S. J. Karrila, *Microhydrodynamics: Principles and Selected Applications* (Courier Corporation, 2013).
- [33] C. Chen, A. I. Packman, and J.-F. Gaillard, Pore-scale analysis of permeability reduction resulting from colloid deposition, *Geophysical Research Letters* **35** (2008).

- [34] S. Chen and G. D. Doolen, LATTICE BOLTZMANN METHOD FOR FLUID FLOWS, *Annual Review of Fluid Mechanics* **30**, 329 (1998).
475
- [35] S. Succi, *The Lattice Boltzmann Equation for Fluid Dynamics and Beyond* (Oxford University Press, 2001).
- [36] S. Succi, R. Benzi, and F. Higuera, The lattice Boltzmann equation: A new tool for computational fluid-dynamics, *Physica D: Nonlinear Phenomena* **47**, 219 (1991).
- [37] H. Chen, S. Chen, and W. H. Matthaeus, Recovery of the Navier-Stokes equations using a lattice-gas Boltzmann method, *Physical Review A* **45**, R5339 (1992).
480
- [38] C. Chen, Z. Wang, D. Majeti, N. Vrvilo, T. Warburton, V. Sankar, and G. Li, Optimization of Lattice Boltzmann Simulation With Graphics-Processing-Unit Parallel Computing and the Application in Reservoir Characterization, *SPE Journal* **21**, 1425 (2016).
- [39] D. Grunau, S. Chen, and K. Eggert, A lattice Boltzmann model for multiphase fluid flows, *Physics of Fluids A: Fluid Dynamics* **5**, 2557 (1993).
485
- [40] Y. Han and P. A. Cundall, Lattice Boltzmann modeling of pore-scale fluid flow through idealized porous media, *International Journal for Numerical Methods in Fluids* **67**, 1720 (2011).
- [41] Y.-H. Qian, D. d’Humières, and P. Lallemand, Lattice BGK Models for Navier-Stokes equation, *EPL (Europhysics Letters)* **17**, 479 (1992).
- [42] S. Chen, D. Martinez, and R. Mei, On boundary conditions in lattice Boltzmann methods, *Physics of Fluids* **8**, 2527 (1996).
490
- [43] M. Fan, J. McClure, Y. Han, Z. Li, and C. Chen, Interaction Between Proppant Compaction and Single-/Multiphase Flows in a Hydraulic Fracture, *SPE Journal* **23**, 1290 (2018).
- [44] C. Chen and D. Zhang, Lattice Boltzmann simulation of the rise and dissolution of two-dimensional immiscible droplets, *Physics of Fluids* **21**, 103301 (2009).
- [45] A. Mohamad and A. Kuzmin, A critical evaluation of force term in lattice Boltzmann method, natural convection problem, *International Journal of Heat and Mass Transfer* **53**, 990 (2010).
495
- [46] J. E. McClure, Lattice Boltzmann Methods for Porous Media (LBPM), <https://github.com/OPM/LBPM> (2019).
- [47] J. E. McClure, Z. Li, M. Berrill, and T. Ramstad, The LBPM software package for simulating multiphase flow on digital images of porous rocks, *Computational Geosciences* **25**, 871 (2021).
- [48] O. Ronneberger, P. Fischer, and T. Brox, U-net: Convolutional Networks for Biomedical Image Segmentation, in *International Conference on Medical image computing and computer-assisted intervention* (Springer, 2015) pp. 234–241.
500
- [49] S. Ioffe and C. Szegedy, Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift, in *International Conference on Machine Learning* (PMLR, 2015) pp. 448–456.
- [50] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, *et al.*,
505 PyTorch: An Imperative Style, High-Performance Deep Learning Library, in *Advances in Neural Information Processing Systems* (2019) pp. 8026–8037.
- [51] D. Kingma and J. Ba, Adam: A Method for Stochastic Optimization, in *Proceedings of the International Conference on Learning Representations* (2015).
- [52] R. Guo, L. E. Dalton, M. Fan, J. McClure, L. Zeng, D. Crandall, and C. Chen, The role of the spatial heterogeneity and
510 correlation length of surface wettability on two-phase flow in a CO₂-water-rock system, *Advances in Water Resources* **146**, 103763 (2020).
- [53] L. E. Dalton, D. Tapriyal, D. Crandall, A. Goodman, F. Shi, and F. Haeri, Contact Angle Measurements Using Sessile Drop and Micro-CT Data from Six Sandstones, *Transport in Porous Media* **133**, 71 (2020).
- [54] Z. Li, J. E. McClure, and T. Ramstad, Bentheimer Sandstone Two-Fluid Flow Simulation Resembling Special Core Analysis Protocol, https://www.digitalrocksportal.org/projects/326/analysis_data/465/ (2020).
515
- [55] B. Zhao, C. W. MacMinn, B. K. Primkulov, Y. Chen, A. J. Valocchi, J. Zhao, Q. Kang, K. Bruning, J. E. McClure, C. T.

Miller, *et al.*, Comprehensive comparison of pore-scale models for multiphase flow in porous media, *Proceedings of the National Academy of Sciences* **116**, 13799 (2019).

520 [56] A. T. Mohan, N. Lubbers, D. Livescu, and M. Chertkov, Embedding Hard Physical Constraints in Convolutional Neural Networks for 3D Turbulence, in *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations* (2020).

[57] X.-H. Zhou, J. E. McClure, C. Chen, and H. Xiao, Predicting flow fields in porous media with neural networks, <https://github.com/xuhuizhou-vt/Flow-field-prediction-in-porous-media> (2022).